REMARKS

The Examiner is thanked for withdrawing the prior objection to the drawings, as well as the prior rejection under 35 USC §112, second paragraph.

The subject matter added to claim 1 (the "given timeout period") is found at page 11, lines 4-6 (see, e.g., step 214 in Figure 2). The subject matter added to claim 9 includes "first" and "second" timeout periods. With respect to the "first timeout period" of claim 9, see page 10, lines 29-31 (e.g., step 208) and page 12, lines 12, lines 18-22 (e.g., steps 306-308); with respect to the "second timeout period" of claim 9, see page 11, lines 4-6 (step 214). (Thus, the "second timeout period" in claim 9 is another way of referring to the "given timeout period" as recited in claim 1). The subject matter of new claim 18 is found at page 12, lines 21-22.

Claims 1 and 5-8 are rejected under 35 USC 103(a) as being unpatentable over Jeffords, U.S. Publication No. 20010042139 in view of Janis, U.S. Patent No. 5,263,165, further in view of newly-cited U.S. Patent No. 6,662,198 to Satyanarayanan et al.

Claims 9-10 and 13-17 are rejected under 35 USC 103(a) as being unpatentable over Jeffords, Janis and Satyanarayan as applied above, further in view of Phillips, U.S. Patent No. 7,058,696.

The Examiner is requested to reconsider these rejections in view of the claim amendments.

In rejecting claims under 35 USC §103(a), the Examiner bears the initial burden of presenting a *prima facie* case of obviousness. *In re Rijckaert*, 9 F.3d 1531, 1532 (Fed. Cir.1993). A "patent [claim] composed of several elements is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art." *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 127 S. Ct. 1727, 1741 (2007). The Federal Circuit has also cautioned not to "allow hindsight reconstruction of references to reach the claimed invention without any explanation as to how or why the references would be combined to produce the claimed invention." *Innogenetics, N.V. v. Abbott Labs.*, 512 F.3d 1363, 1373 n.3 (Fed.Cir.2008).

The "test for obviousness is what the combined teachings of the references would have suggested to one of ordinary skill in the art." *In re Young*, 927 F.2d 588, 591 (Fed.Cir.1991).

Jeffords describes a method and apparatus for accessing resource objects contained in a distributed memory space in a communications network. A "resource object" is a software object whose data may be important [0223]. In this system, the distributed memory space is divided into a plurality of memory pools, with each pool containing a collection of resource objects [0221]. A plurality of resource manager objects are also provided, with each resource manager object having an associated set of memory pools and a registry of network unique identifiers for the resource objects in those pools. A resource manager provides a registry of resource objects and a mechanism for communicating with other resource managers [0217]. The patent application also describes that the system affords a so-called "relativistic view" of state of a plurality of objects, where each object generates a state vector representing that object's view of its own state and the state of all other objects [0013]. Using a "state vector exchange protocol," each object sends its state vector to other objects, and each object maintains a state matrix of the state vectors [0057].

Figure 4 in Jeffords illustrates a state vector generated by an object A. In this example, object A thinks itself to be in state 1 and thinks object B is in state n, object C is in state 3, and object D is in state 2 [0180]. Figure 5 is a representative state matrix generated by the combination of state vectors for objects A through D. In this case, each object has provided a state vector (row) to the matrix. The column vector for object A indicates that object A thinks itself to be in state 1, object B thinks object A is in state 3, object C thinks object A is in state n, and object D thinks object A is in state 3 [0185].

In Jeffords, it is assumed that a distributed application runs with one or more processes at any given time. On start up, each process initializes an internal contact status matrix (CSM) and attempts to contact a set of other processes. Figure 8 illustrates a method for updating state information in this matrix. The method begins by having a particular process (process 1) update its own contact status vector. Next, the process 1 attempts to contact the other systems. If a contact to another system is successful, the process 1 updates its own state vector to provide this indication. The process 1 then sends

out a copy of its state vector to the other systems, e.g., processes 2-n (107). Then, process 1 continues to contact other systems (108); when all processes have been attempted to be contacted, the method ends. Using the discovery process, each process learns about the relative state of all other processes in the network. In particular, and as described at [0057], this "is done through a resource manager state vector exchange protocol, i.e., the use of state matrices filled by the vectors received at each resource manager … Once a stable (deterministic) network state has been achieved (all active … processes are reporting the same state information about each other), synchronization of resource pools is initiated."

In Jeffords, and as described in the flowchart of Figure 8, an individual process attempts to contact other processes (that will be required for use by the distributed application), and that process updates its own state vector accordingly before transmitting that updated vector to the other processes. In Jeffords, and with all due respect to the Examiner's interpretation, a process does not receive a state vector passed from some other process which it then modifies - e.g., by writing into the received vector its own state information. Moreover, in Jeffords, running the state vector exchange protocol appears to be for the purpose of ensuring that the active processes know when they are reporting the same state information about each other; this enables the system there to initiate operation of the distributed application that will use these processes. The protocol is not, as in the subject disclosure (and, again, with respect to the Examiner's contrary position), establishing a guarantee that the servers (that have modified the bits in the knowledge bit vector) have reached an agreement that a file submit process can now proceed. Indeed, the distributed application in Jeffords has nothing to do with file submission; rather, Jeffords describes a distributed call management application [0209] or a text-based conferencing application [0210].

Moreover, Jeffords does not teach that his state vector exchange protocol must be carried out within a given timeout period. Rather, the reference appears to teach the opposite: [0052] the "processes that compose an RRM system may be started in a completely asynchronous manner, that is, they may be started in any order and at any time."

Turning to the secondary references, Janis is cited for its teaching of "access control on system resources such as files within a distributed data processing system having multiple resource managers." Satyanarayanan discloses a system and method for asynchronously sharing, backing up, and distributing data based on immutable entries. The reference is cited for its teaching of transmitting in bits and using a bit to indicate file status, such as the archive status of a backup. Philips is cited for its teaching of "encoding given information about [a] file into a temporary identifier and having the identifier output to each of a set of servers to which a server has connectivity." The Examiner's reading of each secondary reference, in of itself, is not contested. Moreover, because "[n]on-obviousness cannot be established by attacking references individually where the rejection is based upon the teachings of a combination of references," *In re Merck & Co.*, 800 F.2d 1091, 1097 (Fed.Cir.1986), the undersigned will not attack the combination(s). Rather, the issue here is the deficiency in the primary reference Jeffords.

As noted above, the patent does not address file submissions across a distributed computer network, nor does it disclose any type of distributed agreement protocol involving the modifying of a received bit vector, let alone such a protocol using vector exchange where the exchange must be carried out within a given timeout period.

Further, with respect to claim 9, there are now two (2) distinct timeout periods, a "first" timeout during which the "file push" (which is not claimed in claim 1) must complete, and a "second" timeout during which the vector agreement protocol must complete. No reference of record describes these features, thus, the combination "as a whole" cannot either.

Thus, Jeffords/Janis/Satyanarayan or Jeffords/Janis/Satyanarayan/Phillips still lacks at least the following limitations now positively recited in the claims:

Claim 1:

"in response to receipt of a submission of a file at a given server, accepting the submission at the given server only if a given subset of the set of distributed servers reach an agreement to the submission within a given timeout period, where the agreement is determined using a data exchange protocol that includes sub-steps as follows:

passing a bit vector from a first server to a second server, the bit vector including a first indication that the first server has knowledge of the file;

upon receipt of the bit vector at the second server, having the second server modify the bit vector to include, together with the first indication, a second indication that the second server also has knowledge of the file;

having the second server pass the bit vector, which includes the first and second indications, to one or more other servers in the given subset; and

upon a given state being reached, as indicated by at least the first and second indications in the bit vector, determining that the agreement has been reached within a given timeout period;"

Claim 9:

"if the file has been successfully pushed to each of the other servers within a first timeout period, having the given server initiate a data exchange protocol to each of the other servers to which the given server has connectivity, where the data exchange protocol includes sub-steps as follows:

passing a knowledge bit vector among the given server and the other servers;

having each server that receives the knowledge bit vector modify the knowledge bit vector to indicate that server's knowledge of the file; and

based on the knowledge bit vector as modified, determining whether a quorum of the servers have reached a given state within a second timeout period; and

if the quorum of servers reach the given state within the second timeout period, accepting the file for submission."

Claim 18:

"wherein the first timeout period is a function of a size of the file."

For these reasons, a Notice of Allowance is respectfully requested.

Respectfully submitted,

/David H. Judson/

By: _____

David H. Judson, Reg. No. 30,467

ATTORNEYS FOR APPLICANT

February 19, 2009